Resource Allocation in Egalitarian Agent Societies

Ulle Endriss[1], Nicolas Maudet[2], Fariba Sadri[1] and Francesca Toni[1]

[1] Department of Computing, Imperial College London
Email: {ue,fs,ft}@doc.ic.ac.uk

[2] School of Informatics, City University, London
Email: maudet@soi.city.ac.uk

Talk Overview

- Resource allocation by negotiation in multiagent systems
  definition of our negotiation framework
- Measuring social welfare in egalitarian societies
  what are social welfare functions? and why egalitarianism?
- Acceptability criteria
  what kinds of deals should an "egalitarian" agent accept?
- Emergence of global effects from local actions
  sufficiency and necessity of certain deals for optimal outcomes
- Conclusion
  summary and future work

Resource Allocation by Negotiation

- Finite set of *agents* $\mathcal{A}$ and finite set of *resources* $\mathcal{R}$.
- An *allocation* A is a partitioning of $\mathcal{R}$ amongst the agents in $\mathcal{A}$.
  Example: $A(i) = \{r_3, r_7\}$ — agent i owns resources r_3 and r_7
- Agents may engage in negotiation to exchange resources in order to benefit either themselves or society as a whole.
- A *deal* $\delta = (A, A')$ is a pair of allocations (before/after).

Utility and Social Welfare

- Every agent $i \in \mathcal{A}$ has a *utility function* $u_i : 2^{\mathcal{R}} \to \mathbb{R}$.
  Example: $u_i(A) = u_i(A(i)) = 577.8$ — agent i is pretty happy
- A *social welfare ordering* formalises the notion of a society's "preferences" given the preferences of its members (the agents).
  Example: the *utilitarian* social welfare function sw_u:

$$sw_u(A) \quad = \quad \sum_{i \in \mathcal{A}} u_i(A)$$

Egalitarian Social Welfare

The first objective of an *egalitarian* society should be to maximise the welfare of its weakest member.

▶ This motivates the *egalitarian social welfare function* sw_e:

$$sw_e(A) \quad = \quad min\{u_i(A) \mid i \in \mathcal{A}\}$$

Allocation A' is strictly preferred over allocation A (by society) iff $sw_e(A) < sw_e(A')$ holds (so-called *maximin-ordering*).

Utilitarianism versus Egalitarianism

- In the multiagent systems literature the utilitarian viewpoint (i.e. social welfare = sum of individual utilities) is usually taken for granted.
- In philosophy/sociology/economics not.
- John Rawls' *"veil of ignorance"* (*A Theory of Justice*, 1971):

  > *Without knowing what your position in society (class, race, sex, ...) will be, what kind of society would you choose to live in?*

- Reformulating the *veil of ignorance* for multiagent systems:

  > *If you were to send a software agent into an artificial society to negotiate on your behalf, what would you consider acceptable principles for that society to operate by?*

- *Conclusion:* worthwhile to investigate egalitarian principles also in the context of multiagent systems.

Acceptability Criteria

An agent i may or may not accept a particular deal $\delta = (A, A')$. Here are some examples for possible *acceptability criteria:*

selfish agent	$u_i(A) < u_i(A')$
selfish but cooperative agent	$u_i(A) \leq u_i(A')$
selfish and demanding agent	$u_i(A) + 10 < u_i(A')$
masochist	$u_i(A) > u_i(A')$

disciple of agent *guru*	$u_{guru}(A) < u_{guru}(A')$
team worker (for team T)	$\sum_{j \in T} u_j(A) < \sum_{j \in T} u_j(A')$

Example for a Protocol Restriction

no more than two agents to be involved in any one deal	$\vert\mathcal{A}^\delta\vert \leq 2$ where $\mathcal{A}^\delta = \{i \in \mathcal{A} \mid A(i) \neq A'(i)\}$

Pigou-Dalton Transfers

A criterion for agents that want to *reduce inequality* ...

In our framework, a *Pigou-Dalton transfer* (between agents i and j) can be defined as a deal $\delta = (A, A')$ with the following properties:

(1) $\mathcal{A}^\delta = \{i, j\}$
(only i and j are involved in the deal)

(2) $u_i(A) + u_j(A) = u_i(A') + u_j(A')$ [could be relaxed to $\leq$]
(the deal is mean-preserving, i.e. overall utility is not affected)

(3) $|u_i(A') - u_j(A')| < |u_i(A) - u_j(A)|$
(the deal reduces inequality)

Pigou-Dalton transfers capture certain egalitarian principles; but are they sufficient as acceptability criteria to guarantee optimal outcomes of negotiations for society?

Example

Consider the resource allocation problem with $\mathcal{A} = \{bob, mary\}$, $\mathcal{R} = \{glass, wine\}$, and initial allocation A:

$A(bob) = \{glass\}$	$A(mary) = \{wine\}$
$u_{bob}(\{\,\}) = 0$	$u_{mary}(\{\,\}) = 0$
$u_{bob}(\{glass\}) = 3$	$u_{mary}(\{glass\}) = 5$
$u_{bob}(\{wine\}) = 12$	$u_{mary}(\{wine\}) = 7$
$u_{bob}(\{glass, wine\}) = 15$	$u_{mary}(\{glass, wine\}) = 17$

The "inequality index" for allocation A is 4 (minimal!).

But allocation A' with $A'(bob) = \{wine\}$ and $A'(mary) = \{glass\}$ would result in higher egalitarian social welfare (5 instead of 3).

Hence, Pigou-Dalton deals alone are not sufficient to guarantee optimal outcomes (they also don't cover deals between more than two agents).

▶ We need a more general acceptability criterion.

Equitable Deals

- Let $\delta = (A, A')$ be a deal.
- $\mathcal{A}^\delta = \{i \in \mathcal{A} \mid A(i) \neq A'(i)\}$ is the set of agents involved in δ.
- We call δ *equitable* iff the following holds:

$$min\{u_i(A) \mid i \in \mathcal{A}^\delta\} \quad < \quad min\{u_i(A') \mid i \in \mathcal{A}^\delta\}$$

(Intuitively, this is egalitarianism "at the local level".)

Maximin-rise implies Equitability

A first connection between our "global" and "local" measures:

Lemma 1 If $sw_e(A) < sw_e(A')$ then $\delta = (A, A')$ is equitable.

Proof. Because any deal that improves social welfare must involve the (previously) poorest agent(s) and increase its (their) utility.

What about Global Effects of Local Actions?

Note that the converse of Lemma 1 does not hold!
Example: any equitable deal only involving the very richest agents

▶ To be able to always detect the effects of equitable deals at the society level we need a finer measure of social welfare.

The Leximin-ordering

Every allocation A gives rise to an *ordered utility vector* $\vec{u}(A)$: compute $u_i(A)$ for all $i \in \mathcal{A}$ and present results in increasing order.

Example: $\vec{u}(A) = \langle 0, 5, 20 \rangle$ means that the weakest agent enjoys utility 0, the strongest utility 20, and the middle one utility 5.

The *leximin-ordering* $\prec$ over allocations is defined as follows:

$$A \prec A' \quad \text{iff} \quad \vec{u}(A) \text{ lexically precedes } \vec{u}(A')$$

Example: $A \prec A'$ for $\vec{u}(A) = \langle 0, 6, 20, 29 \rangle$ and $\vec{u}(A') = \langle 0, 6, 24, 25 \rangle$

Equitability implies Leximin-rise

Lemma 2 If $\delta = (A, A')$ is equitable then $A \prec A'$.

Proof. [see paper]

Termination

Lemma 3 (Termination) There can be no infinite sequence of equitable deals, i.e. negotiation will always terminate.

Proof. The space of distinct allocations is finite and, by Lemma 2, every equitable deal results in a strict rise wrt. the leximin-ordering.

Guaranteed Optimal Outcomes

Theorem 1 (Sufficiency) Any sequence of equitable deals will eventually result in an allocation with maximal social welfare.

Proof. By Lemma 3, negotiation must terminate. Assume the final allocation A is *not* optimal, i.e. there exists an allocation A' with $sw_e(A) < sw_e(A')$. But then, by Lemma 1, the deal $\delta = (A, A')$ would be equitable (contradicts assumption that A is final).

Discussion

▶ Note that <u>*any*</u> sequence of (equitable) deals will eventually result in an optimal allocation.

▶ Agents can act *locally* and do not need to be aware of the global picture (the positive global effect is guaranteed by the theorem).

Necessity of Complex Deals

Theorem 2 (Necessity) For every deal δ, there is an instance of the resource allocation problem (utility functions and initial allocation) such that no sequence of equitable deals excluding δ could result in an allocation with maximal social welfare.

Proof. [by construction; see paper]

Discussion

▶ Very complex deals (involving any number of agents and resources) may be necessary to guarantee optimal outcomes.

Conclusion and Future Work

- Egalitarian social welfare is relevant to multiagent systems.
- *Welfare engineering:* Force a desired behaviour at society level by engineering a suitable negotiation policy for individuals.
- Other examples:
  - AAMAS-2003: *utilitarian* social welfare and *selfish* agents [see also related work by T. Sandholm on task allocation]
  - *Elitist societies:* social welfare depends on the happiest agent (agents cooperate to support their "champion" to make sure at least one of them achieves their goal)
- Maybe certain *types* of deals (say, involving only two agents) can guarantee optimal outcomes for restricted domains? [some results for the utilitarian case, but not for the egalitarian]
- Develop *protocols* for multi-item/multi-agent trading.